Mon. Not. R. Astron. Soc. **000**, 000–000 (0000) Printed 29 September 2014 (MN LATEX style file v2.2)

The variability of the Crab Nebula in radio: No radio counterpart to gamma-ray flares

M. F. Bietenholz[1,2], Y. Yuan[3], R. Buehler[4], A. P. Lobanov[5,6], and R. Blandford[3]
[1] *Hartebeesthoek Radio Observatory, PO Box 443, Krugersdorp, 1740, South Africa*
[2] *Department of Physics and Astronomy, York University, Toronto, M3J 1P3, Ontario, Canada*
[3] *Kavli Institute for Particle Astrophysics and Cosmology, Stanford University, Menlo Park, CA 94025, US*
[4] *Deutsches Elektronen Synchrotron (DESY), Platanenalee 6, 15738, Zeuthern, Germany*
[5] *Max-Planck-Institut für Radioastronomie, Auf dem Hügel 69, 53121 Bonn, Germany*
[6] *Institut für Experimentalphysik, Universtität Hamburg, Luruper Chaussee 149, 22761, Hamburg, Germany*

arXiv:1409.7627v1 [astro-ph.SR] 26 Sep 2014

29 September 2014; *Accepted to MNRAS*

ABSTRACT

We present new Jansky Very Large Array (VLA) radio images of the Crab Nebula at 5.5 GHz, taken at two epochs separated by 6 days about two months after a gamma-ray flare in 2012 July. We find no significant change in the Crab's radio emission localized to a region of < 2 light-months in radius, either over the 6-day interval between our present observations or between the present observations and ones from 2001. Any radio counterpart to the flare has a radio luminosity of $\lesssim 2 \times 10^{-4}$ times that of the nebula. Comparing our images to one from 2001, we do however find changes in radio brightness, up to 10% in amplitude, which occur on decade timescales throughout the nebula. The morphology of the changes is complex suggesting both filamentary and knotty structures. The variability is stronger, and the timescales likely somewhat shorter, nearer the centre of the nebula. We further find that even with the excellent *u-v* coverage and signal-to-noise of the VLA, deconvolution errors are much larger than the noise, being up to 1.2% of peak brightness of the nebula in this particular case.

Key words: supernova remnants

1 INTRODUCTION

The Crab Nebula is one of the most intensely studied objects in astrophysics, yet it retains the power to surprise us. It is the remnant of a supernova explosion in the year AD 1054, which was witnessed by Chinese and other astronomers. The presently visible remnant is bright at all observable wavelengths. It has, in particular, long been observed in the radio, where it is one of the brightest sources at GHz frequencies (see, e.g., Hogg et al. 1969; Wilson 1972, and references therein). It also contains one of the first known pulsars, PSR B0531+21. The Crab Nebula is the prototype of a pulsar-powered nebula, commonly known as a pulsar wind nebula (PWN), where the rotational energy lost by the pulsar as it spins down powers presently visible nebula (see Bühler & Blandford 2014; Hester 2008, for recent reviews). The energy input from the pulsar, which emerges in the form of a wind of magnetic field and relativistic particles has inflated a bubble of relativistic fluid. Since the pulsar spindown is quite regular, the rate of this energy input is quite steady, although it decreases slowly, in the case of the Crab at $\sim$0.14% yr^{-1}.

A reason to re-observe the Crab in the radio as well as at other wavelengths occurred when it was recently discovered that the Crab produces substantial flares in gamma-ray emission. Several such flares have now been observed, where the emission at photon energies > 100 MeV increases by more than a factor of two on timescales of days (e.g. Buehler et al. 2012; Striani et al. 2011; Tavani et al. 2011). These flares are not yet well understood — see discussions in e.g. Arons (2012); Bykov et al. (2012); Lyutikov, Balsara & Matthews (2012); Lyutikov (2014); Sturrock & Aschwanden (2012), and Komissarov & Lyutikov (2011), although they likely involve regions moving relativistically towards us and/or perhaps magnetic reconnection. While there are occasionally small sudden changes in the spindown rate of the pulsar known as glitches, Espinoza et al. (2010) and Mickaliger et al. (2012) have shown glitches do not seem to occur in conjunction with the gamma ray flares.

On 2012 July 3, the Fermi Large Area Telescope (Fermi-LAT) detected a new flare from the Crab (Ojha et al. 2012) in gamma-rays, with the flux at ener-

© 0000 RAS

gies > 100 MeV having increased to $(5.5 \pm 0.7) \times 10^{-6}$ photons cm^{-2} s^{-1} from levels near its long-term average level of 2.75×10^{-6} photons cm^{-2} s^{-1} (Nolan et al. 2012). The daily average flux reached a peak of $(6.2 \pm 0.8) \times 10^{-6}$ photons cm^{-2} s^{-1}, over twice the long-term average value, but by July 8 it had decayed back to the long-term average level (preliminary values from our reduction, Buehler). Note that the above flux values are the totals, consisting of the combined flux from both the pulsar and the nebula. The average flux from the nebula is only 0.6×10^{-6} photons cm^{-2} s^{-1}.

The short timescales of the flares imply that the emitting region is quite small: unless there is ultrarelativistic beaming, the flare emission must come from a region of $\lesssim 3 \times 10^{-4}$ pc, corresponding to $\lesssim 0\farcs03$ in size (Abdo et al. 2011; Bühler & Blandford 2014, and references therein). A similar conclusion was reached by modelling the combined gamma-ray, X-ray and radio spectrum (Lobanov, Horns & Muxlow 2011; Meyer, Horns & Zechlin 2010).

Since the flaring component of the emission is not pulsed, it is thought to originate in the nebula rather than the pulsar (Abdo et al. 2011; Buehler et al. 2012). Together with the short timescales, this implies that it must be synchrotron emission, arising from electrons with energies of order 10^{15} eV. In gamma rays, the nebula therefore brightened by almost a factor of 6 during the July 2012 flare. Although this flare was slightly less luminous than the earlier flares, it still represents a very energetic event, as well as only the fourth time such high fluxes had been observed in four years of Fermi observations.

The resolution of Fermi-LAT is on the order of 1°, therefore the gamma-ray emission cannot be localized within the nebula (which is about $7' \times 5'$ in extent). Telescopes from infrared to X-ray observed the Crab during or shortly after previous flares and found no unusual emission (Abdo et al. 2011; Buehler et al. 2012, and references therein).

Several locations in the inner nebula have in fact been suggested as the location of the gamma-ray flares. They include the "inner knot" first identified by Hester et al. (1995)[1], as well as the "anvil" region suggested by Tavani et al. (2011). As of yet, however, the location of the gamma-ray flaring within the nebula has not yet been identified.

Since radio emission is often associated with gamma-ray emission, and since arcsecond resolution is easily achievable in the radio, radio observations seem a natural choice to attempt to pinpoint the location of the flare activity. In more energetic flares, the spectrum of the flaring part of the gamma-ray emission appears to have a cutoff below photon energies of ∼200 MeV, but the nature of the spectrum of the less energetic flares, like the one of 2012 July, is not clear (Bühler & Blandford 2014). Although any extrapolation of the spectral energy distribution seen in the gamma-ray flares to radio frequencies (photon energies of order 10^{-5} eV), seems dangerous, the source of the energy seems likely to be the magnetic field, and one might expect disturbances in the magnetic field to produce changes in the radio-frequency synchrotron emission. Furthermore, the electrons responsible for the flaring gamma-ray emission should certainly still be energetic enough to produce radio emission long after the flare event.

The very energetic electrons producing the gamma-ray emission have very short lifetimes, as they rapidly lose their energy due to synchrotron losses. By contrast, the lifetimes of the less energetic electrons producing radio emission is generally longer than the age of the nebula. Therefore, even though the high-energy emission is short-lived, any response to it in particles emitting at radio frequencies would be adiabatic, and any corresponding radio emission should be discernible until the particles diffuse into the body of the nebula.

In fact, the origin of the relativistic electrons producing the radio synchrotron emission, i.e. those having energies $\lesssim 10$ GeV, is a long standing puzzle. These electrons have synchrotron lifetimes long compared to the age of the nebula. If the radio-emitting electrons are continuously injected into the nebula along with the higher energy ones, the total number of particles is larger than produced by current theories (Arons 2012; Hibschman & Arons 2001). However, if they are not injected by the pulsar, but accelerated by some other process in the nebula, then the existence of radio features, called "wisps", associated with the termination shock (Bietenholz et al. 2004; Bietenholz, Frail & Hester 2001; Bietenholz & Kronberg 1992) is a puzzle, as is the continuity of the spectral energy distribution from the radio through to the optical and above (Aharonian & Atoyan 1998). The first of these issues may not be a problem: Olmi et al. (2014) have shown that even if the radio-emitting electrons are produced in the body of the nebula, their long lifetimes would allow them to diffuse essentially everywhere in the nebula, and they would still be expected to be present near the termination shock, and that variations of the magnetic field near the shock would then naturally produce the radio wisps. Komissarov (2013) has shown that the radio-emitting electrons could plausibly be accelerated by magnetic dissipation in the vicinity of the filaments.

In any case, since the acceleration of the higher energy particles responsible for the gamma-ray flare is likely to happen near the pulsar, the detection of (or upper limit to) any radio counterpart to the flares would shed light on the low-energy tail of the acceleration mechanism responsible for the gamma-ray flares, and might therefore also shed light on the more general process for the production of the radio-emitting electrons.

Radio observations of the Crab in response to gamma-ray flares have been previously undertaken. Firstly, Lobanov, Horns & Muxlow (2011) obtained 1.4-GHz VLBI observations, most sensitive to features $< 0\farcs1$ in size, 44 d after the flare of 2011 Sep. They detected two weak features, called knots C1 and C2, which they tentatively identified as possible radio counterparts to the gamma-ray flares. The flux density of these knots C1 and C2 was 0.5 and 0.4 mJy and their distance from the pulsar $5\farcs4$ and $1\farcs6$, respectively. Secondly, Weisskopf et al. (2013) obtained a series of Karl G. Jansky Very Large Array (VLA), as well as *Chandra* X-ray and Keck optical, observations at $t = 3$ to 89 d following the strong 2011 April flare. Their radio observations had arcsecond resolution. They report no compact radio emission features > 0.5 mJy except for the pulsar, although their radio observations were not sensitive enough to detect C1 or

[1] This feature was called "Knot 1" in Hester et al. (1995)

C2 at the flux densities found by Lobanov, Horns & Muxlow (2011).

In response to the 2012 July gamma-ray flare, we therefore undertook Target of Opportunity observations of the Crab with the VLA to search for radio emission in response to the gamma-ray flare, and confirm the tentative detection of Lobanov, Horns & Muxlow (2011). The 2012 July flare fortuitously occurred while the VLA was in the B-array configuration which is well suited to high-resolution imaging of the Crab, giving arcsecond resolution at 5 GHz.

A confirmed detection of radio emission in response to a gamma-ray flare would localize the flaring region in the complex structure of the nebula. A detection, or even a non-detection with a sensitive upper limit, would provide an important constraint on the broad-band spectral energy distribution of the emitting material.

Many of the current models of the flares (e.g. Lyutikov 2014; Cerutti et al. 2013; Clausen-Brown & Lyutikov 2012) involve reconnection events, which would produce quasi-monoenergetic particle distributions, which would give rise to synchrotron emission with a spectrum $\propto \nu^{1/3}$ below the peak frequency and which would be unobservably faint in the radio. However, the large energy release would almost certainly cause disturbances in the magnetic field, which should be made visible in radio via the synchrotron emission of long-lived radio-emitting electrons not specifically accelerated by the flaring event.

To our knowledge, there have been no calculations of the expected radio brightness changes in response to a gamma-ray flare. We perform such a calculate below (Section 6) and determine that the radio brightness changes may well be too small to be detectable in our observations. However, this calculation is very model-dependent, and was only performed after we had obtained the target-of-opportunity radio observations which are the subject of the present paper. Particularly given the tentative detection of Lobanov, Horns & Muxlow (2011) we consider that our new observations were well motivated.

2 OBSERVATIONS AND DATA REDUCTION

We obtained two sessions of VLA observations (observing code 12A-486) of the Crab with midpoints 2012 August 20.5 and 26.5 UT, spaced 6.0 days apart, or 49 and 55 days after the onset of a gamma-ray flare. The observations employed a bandwidth 2048 MHz around a central frequency of 5503 MHz, spanned a total of 5 hours per session, and, as already mentioned, were carried out with the array in the B configuration, resulting in baseline lengths between $\sim$2 and $\sim$250 Kλ. The flux density scale was calibrated using observations of 3C 48 and 3C 147, using the Perley-Butler (2010) coefficients, and we used QSO J0559+2353 (QSO B0556+238) as a phase-reference source.

Observations in the B array configuration at 5 GHz do not provide sufficient information at low spatial frequencies to reconstruct the entire nebula. In particular, the observations contain no direct information on spatial scales larger than $\sim 50''$. We employ the strategy for obtaining reliable images of the Crab using only a single array configuration devised by Bietenholz, Frail & Hester (2001); Bietenholz et al. (2004), and we repeat a brief description here. The strategy involves employing maximum-entropy deconvolution using a support made from the B, C and D array configuration observations taken in 1987 to 1988 (Bietenholz & Kronberg 1991, 1990). The support was scaled spatially to account for the expansion of the synchrotron nebula of 0.13% yr^{-1} (Bietenholz et al. 1991), and in brightness to account for the long-term decay of the nebula[2] at 0.202% year^{-1}. The use of the same default for both sessions serves to make any differences between our images at different epochs be only those that are demanded by the data.

3 RESULTS

3.1 Radio images of the Crab

We show the radio image of the Crab nebula at 5 GHz on 2012 August 26 in Figure 1. Since the upgraded VLA has much wider bandwidth than it did during the earlier observations of the Crab, this image has both better image fidelity due to better coverage in the Fourier-transform or u-v plane, as well as lower noise, than any previous radio images of the Crab. In Figure 2 we show a detail of the region around PSR B0531+21.

To determine whether there were any rapid changes in the radio images due to the flaring activity, we formed a difference image between our two epochs. Since both epochs were deconvolved using the same default image, they will be biased to be as similar as is allowed by the visibility measurements, thus reducing the likelihood of spurious difference features due for example to deconvolution errors. Before differencing, both images were convolved to a common restoring beam, for which we chose a size of $1.10'' \times 1.00''$ at p.a. -65°, marginally larger than the value fit to the dirty beams of each individual epoch.

An interferometer is a spatial filter, so the VLA in the B-array at our lowest frequency of 4.7 GHz can reproduce spatial scales smaller than[3] $37''$. Accordingly, we high-pass filtered the difference image with a Gaussian filter with $37''$ FWHM to remove any spurious structure at lower spatial scales (note that our use of a common support for the two epochs should have already served to minimize any such differences).

We show the resulting difference image in Figure 3, and a detail of the centre in Figure 4. The rms differences over the body of the Crab nebula were 0.26 mJy beam^{-1}, and the extrema of the difference image were -1.48 and $+1.56$ mJy beam^{-1}, with the area examined being $\sim$74,000 beam areas. Expressed in terms of the peak brightness of the individual images, which was $\sim$25 mJy beam^{-1} at this resolution, the peak-to-peak difference was about $\pm 6\%$ while the rms value over the body of the nebula was 1.2%.

The overall rippling in the difference image, as we will show in more detail in § 4 below, is mostly due to deconvolution errors and likely does not represent any actual pattern

[2] We take a weighted average of the decay-rate values from Weiland et al. (2011) at 22 and 30 GHz, Aller & Reynolds (1985) at 5 GHz and Vinyaikin (2007) at 927 MHz; our results are not sensitive to small changes in this decay rate.

[3] VLA Observational Status Summary, Table 5; `https://science.nrao.edu/facilities/vla/docs/manuals/oss2014a`

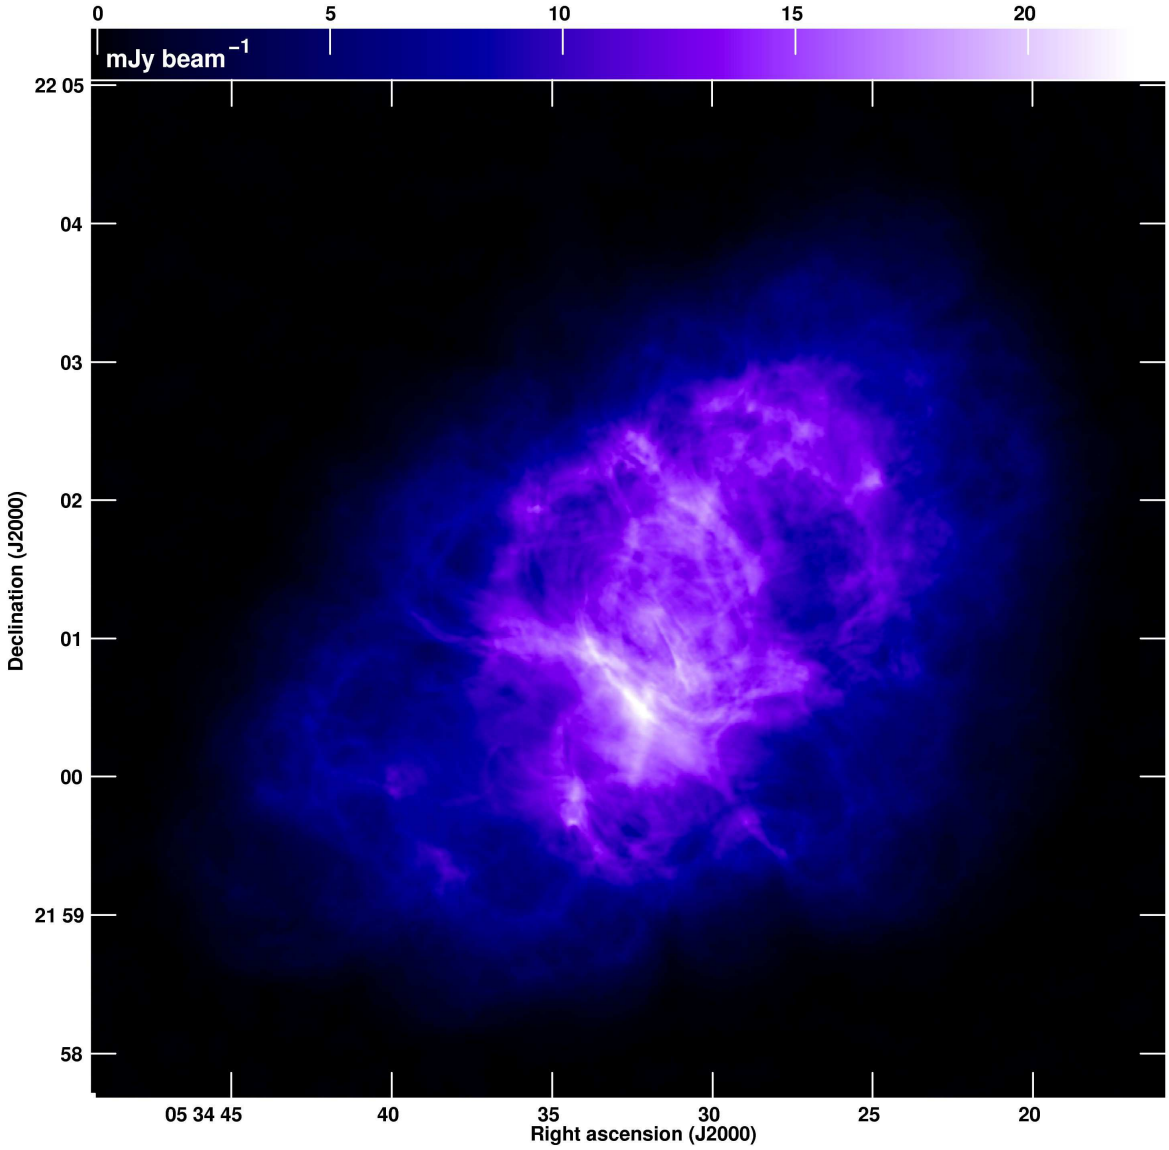

Figure 1. An image of the Crab Nebula at 5.5 GHz on 2012 August 26. The FWHM of the convolving beam was $1.03'' \times 0.99''$ at p.a. $-66°$. The peak brightness was 22.9 mJy beam^{-1}, and the background rms was 28 μJy beam^{-1}. The total flux density recovered in the deconvolution was 592 Jy. Maximum entropy deconvolution was used, with an appropriately scaled default image made from the 1987 and 1988 multi-configuration VLA observations (see text for details). Most of the structure on large spatial scales ($\gtrsim 37''$) is derived from the default image rather than from the present observations. The image has been corrected for the effect of the VLA primary beam response.

of change in the nebula's brightness. No isolated difference feature is visible which might correspond to rapid, localized changes in brightness at the location of the gamma-ray flare event. At the distance of the Crab (2 kpc), a speed of c corresponds to $2\farcs6$ per month. One might therefore expect that any feature related to the sudden energy release of the 2012 July 3 gamma-ray flare would be less than $9\farcs5$ in diameter by the second epoch.

We show a detail of the central region in Figure 4. Note the pattern of elliptical ripples visible around the pulsar. The pattern is very similar to that seen by Bietenholz, Frail & Hester (2001) and Bietenholz et al. (2004), and almost certainly represents real wisp motions over our 6-day interval.

The rms differences over the body of the Crab are $\sim$7 times the rms background value of the difference image. One might therefore conclude that the differences in the two images are due to real changes in the nebula's brightness. As mentioned, however, no particular feature associated with the gamma-ray release is visible. Furthermore, although temporal changes over a large part of the nebula have been seen in difference images spanning larger time intervals in the past, (Bietenholz et al. 2004; Bietenholz, Frail & Hester 2001), the pattern in the present difference images does not resemble those earlier patterns, but rather suggests linear ripples extending over the whole body of the nebula, which seem physically less plausible.

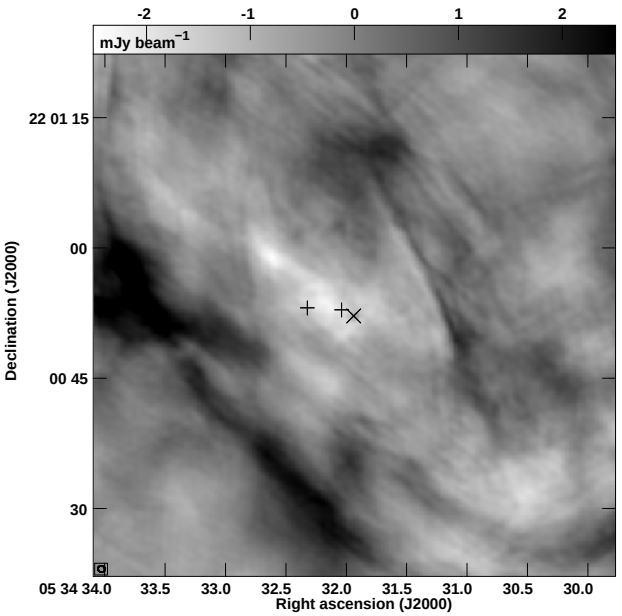

Figure 2. A higher resolution 5.5-GHz image of the central region of the Crab on 2012 Aug. 26. The image is made using the same visibility data as Fig. 1, but using uniform weighting for higher resolution, and multi-resolution CLEAN deconvolution. Since no support information is used, this image is high-pass filtered, and consequently has a mean brightness near 0, but should accurately show the details at small-spatial scales. The FWHM of the restoring beam ($0.80'' \times 0.72''$ at p.a. 68°) is indicated at lower left. We mark the pulsar, PSR B0531+21, with "×" and the knots C1 and C2 from Lobanov, Horns & Muxlow (2011, see text § 6) with "+". The pulsar can be faintly discerned and has a flux density of 1.2 ± 0.3 mJy.

3.2 Central Point Source

A faint unresolved source can be seen in the centre of our images (Figs. 1, 2). It has a 5.5-GHz flux density of 1.05 ± 0.13 mJy on Aug. 20 and 1.17 ± 0.16 mJy on Aug. 26, where the uncertainty includes a contribution due to the difficulty in reliably separating the point source from the background.

We determined the position of this source on images made before any self-calibration (since self-calibration can cause spurious position shifts), and found that it was within 25 ± 40 milli-arcsecond (to the northwest) of the position of the pulsar given by Lobanov, Horns & Muxlow (2011). We conclude that the central point source is coincident with the pulsar to within $0\farcs14$ (3σ upper limit). Its flux density rose by a not statistically significant amount of $(0.12 \pm 0.21$ mJy or 0.12 mJy or (11 ± 19)% over our six-day interval.

4 A CAUTIONARY TALE

The differences between our two epochs of Crab radio observations were about 50× larger than the off-source image rms. Are these differences real? If so, what is their origin? The difference image in Figure 3 suggests striations extending over the whole body of the nebula, but concentrated at particular position angles and spatial frequencies, in other words, originating in small regions in the Fourier transform or u-v-plane. Such a pattern is suggestive of errors due to un-sampled regions of the u-v plane rather than real changes in the nebula's brightness.

Our difference image consists of the difference of two images, each of which was made by Fourier-transforming and then deconvolving an incompletely sampled set of visibility measurements. The deconvolution process attempts to interpolate into the un-sampled regions of the u-v plane in a manner consistent with the known image-plane constraints. In the case of our maximum entropy deconvolution, these constraints are firstly that the true brightness must be positive, and secondly that it be confined to some region in the original image (often termed the "CLEAN" window even in cases where CLEAN is not the deconvolution algorithm used. Note that the images reproduced in Figs. 1 and 3 show only this CLEAN-window sub-region of the complete images, which latter spanned $18'$ in each coordinate). In u-v plane, then, any region which was not sampled is more weakly constrained than those regions which were sampled. In the Fourier transform of the *difference* image, therefore, it is those regions which were not sampled in *both* of our two observing epochs which are more weakly constrained.

We formed the fast Fourier transform (FFT) of the difference image. It is shown in the top panel of Figure 5. For each individual epoch, the beam is the Fourier transform of the sampling function. The *product* of the Fourier transforms of the beams therefore show the regions which were not sampled by both epochs, since in such un-sampled regions, the product is near zero. We show this product of the u-v plane sampling functions, in other words the combined sampling function, in the lower panel of Fig. 5.

It is evident by comparing the two panels of Fig. 5 that the power in the difference image lies predominately at those spatial frequencies where the combined sampling function was zero. We can therefore conclude that the majority of the structure seen in the difference image, despite being many times larger than the background rms brightness, is not real, but due only to errors in the deconvolution, which is necessitated by the incomplete sampling in the u-v plane, or as they are commonly termed "deconvolution errors".

The rms brightness of the difference image over the body of the Crab is 306 μJy beam^{-1} (at a resolution of $1.10'' \times 1.00''$). Expressed in terms of the peak brightness, the difference rms is 1.2%, while the difference extrema are ∼7%. We conclude that the deconvolution errors, even in this case of the excellent u-v-coverage afforded by the VLA wide-band system and several hours of observations, are at an rms level of 1.2% of the image brightness, with extrema at 7%. The off-source rms in these images was ∼30 μJy beam^{-1} so the deconvolution errors are larger than the off-source noise level by an order of magnitude! We made a similar difference image using CLEAN rather than maximum entropy deconvolution, and found that the deconvolution errors were larger by ∼30% than they were for the maximum entropy deconvolution. CLEAN is known to perform more poorly on extended sources (Sault & Oosterloo 1996; Narayan & Nityananda 1986), so the fact that CLEAN deconvolution errors are higher than those obtained from maximum entropy is expected for an extended object like the Crab.

The effective uncertainty in the brightness of the image is thus much larger than the noise, and the image is strongly dynamic range limited. We note that the fractional errors

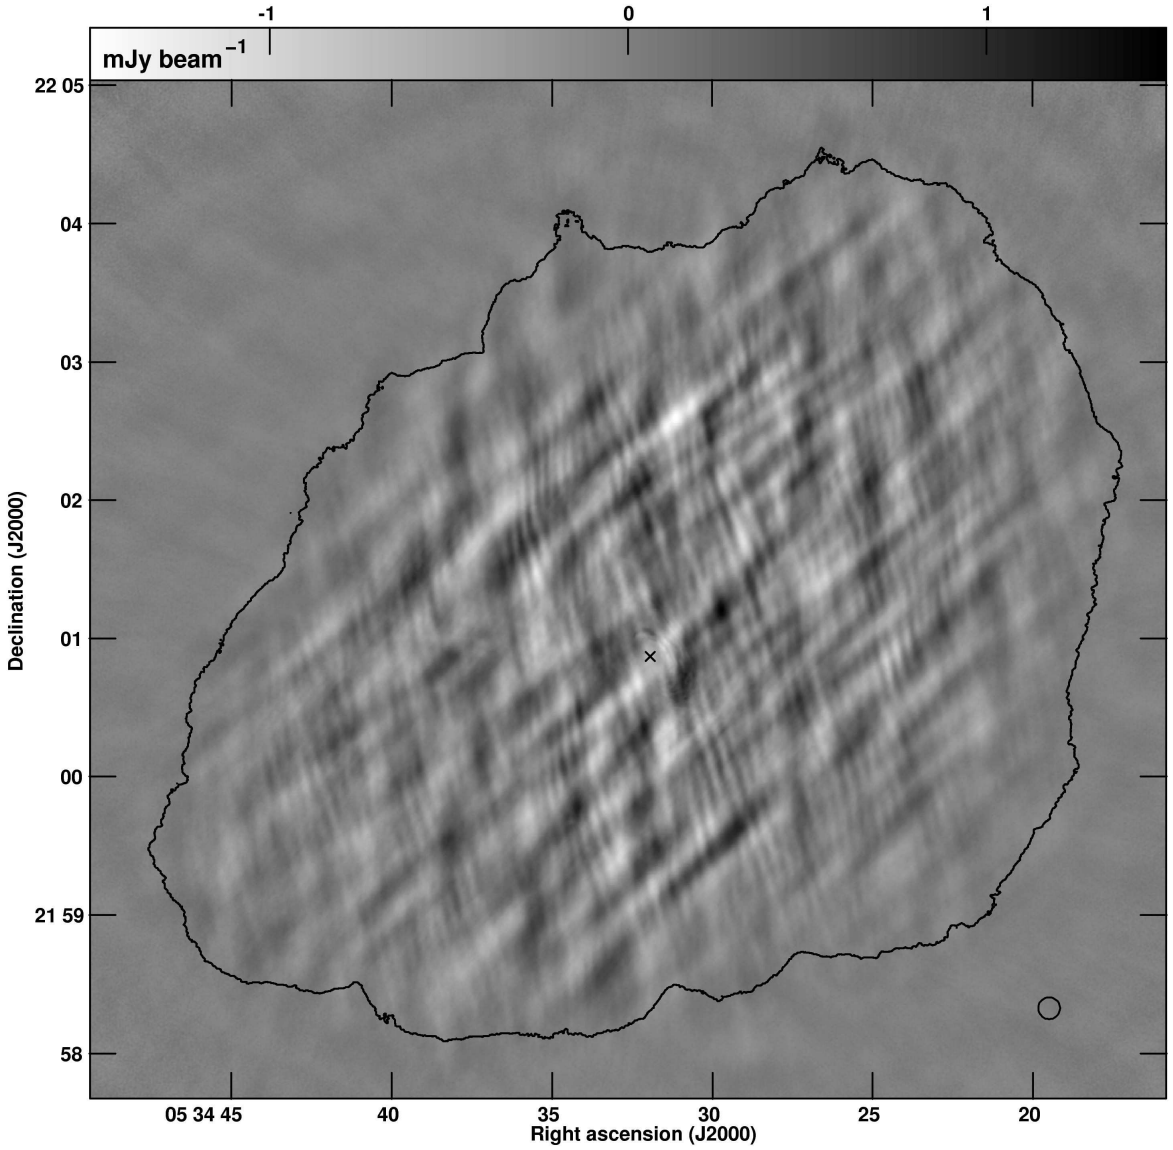

Figure 3. A difference image of the Crab, showing the changes between 2012 August 20 and 26 (with the former subtracted from the latter). The white "×" in the centre gives the position of PSR B0531+21 (Lobanov, Horns & Muxlow 2011). For reference, we also show the 1% contour of the August 26 image. The difference image has been high-pass filtered at $37''$, since the observations were not sensitive to structure on larger scales. Before forming the difference, both images were convolved to a common restoring beam size of $1.10'' \times 1.00''$ at p.a. $-65°$. The off-source rms in the difference image in the region around the Crab was ~ 41 μJy beam^{-1}. No isolated feature is visible which might correspond to the location of the gamma-ray flare event. The circle at lower right shows a diameter of $9\farcs5$, which is the diameter that a feature expanding isotropically with speed c since the gamma-ray flare onset on 2012 July 3 would have at distance 2 kpc.

are probably smaller in the case of spatially smaller sources, and for single unresolved sources the image error will likely approach the thermal noise.

5 DIFFERENCES OVER LONGER TIMESCALES

On the one-week timescale between our two 2012 images, we found above that differences were small and only in the centre of the nebula were convincing changes corresponding to rapid motions seen. However, slower motions are known to occur both in the centre and over the body of the nebula (Bietenholz et al. 2004; Bietenholz, Frail & Hester 2001; Bietenholz & Kronberg 1992). We compare our new images to those made using VLA observations in 2001, also at 5 GHz and using the B array configuration. The 2001 observations are described in Bietenholz et al. (2004).

One goal of this comparison is to measure the expansion rate of the nebula. We defer discussion of the measured expansion rate to a future paper. In order to avoid biases introduced by scaling the default in the deconvolution with the expansion rate measured by Bietenholz et al. (1991), we image using only the B array configuration data without us-

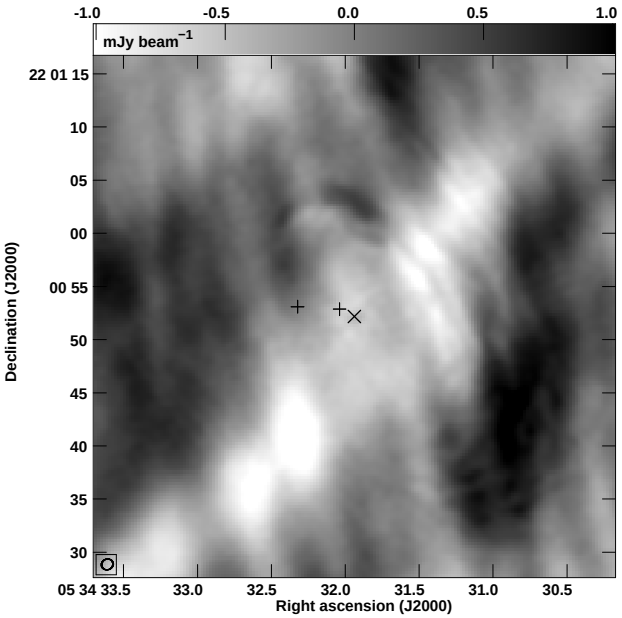

Figure 4. A detail from Fig. 3, showing the region near the pulsar in the Aug. 26 − Aug. 20 difference image. The FWHM of the restoring beam ($1.10'' \times 1.00''$ at p.a. $-65°$) is indicated at lower left. We mark the pulsar with "×" and the knots C1 and C2 from Lobanov, Horns & Muxlow (2011, see text § 6) with "+". The pulsar is faintly visible as a positive excursion, implying that the pulse-averaged flux density of the pulsar increased by ∼200 μJy beam^{-1} between the two dates. A series of ripples due to the outward motion of the elliptical wisps is also visible.

ing any default or support. As a consequence, our images are hi-pass filtered, and may include genuinely negative values. Since maximum entropy deconvolution enforces positivity and therefore cannot deal with true negative values in the images, we turn to CLEAN deconvolution. To avoid spurious differences due to differing u-v coverage either at the short or long end of the spacings covered by the B configuration, we further hi-pass filtered both images with a Gaussian of FWHM $20''$, and also convolved both to a common restoring beam of $2.00'' \times 1.80''$ at p.a. $80°$.

The flux scaling and registration in RA and Dec. might be slightly discrepant between our two images, which were made using self-calibration. Such discrepancies would have a significant effect on difference images. We therefore used the Miriad program IMDIFF, which calculates how to make one image most closely resemble another by calculating unbiased estimators for the scaling in size, e, the scaling and the offset in flux density, A and b respectively, and the offsets in RA and Dec, x and y respectively, needed to make the second image most closely resemble the first.

Finally we show in Figure 6 the *residual* image after finding the best fit values of e, A, b, x and y above We treat e here merely as a nuisance parameter and defer the discuss the expansion rate to a future paper. We discuss here the slow but significant changes in the radio brightness over the whole body of the nebula not attributable to the overall expansion. This difference image shows the changes that have occurred in the Nebula over the 11.4 yr period between 2012 and 2001 (at spatial scales between $\sim 2''$ and $20''$).

The changes in the brightness over the 11.4 yr interval between 2001 and 2012 involve a complex series of ripples. The extrema in the difference image are −7.50 and +8.34 mJy beam^{-1}. The rms over the nebula is 1.20 mJy beam^{-1}, while the rms in a region $1'$ in radius around the pulsar is 1.97 mJy beam^{-1} (FWHM beam area = 4.08 arcsec^{-2}). By comparison the rms variation between the two 2012 epochs (Fig. 3) were four times smaller with rms 0.30 mJy beam^{-1} over the nebula (albeit with a smaller FWHM beam area of only 1.25 arcsec^{-2}). Also, comparing the pattern of differences the pattern visible of the 11.4-year interval are much less regular, and a FFT of the 11.4-year difference image shows no patterns related to the u-v coverage. We conclude therefore that most of the changes visible in Fig. 6 are real, although this difference image will also contain artifacts due to deconvolution errors.

6 DISCUSSION

In response to a gamma-ray flare in the Crab nebula in 2012 July, we observed the Crab Nebula using the VLA with $\sim 1''$ resolution (B-array configuration) in the hopes of seeing a radio counterpart to the gamma-ray flare which would allow us to accurately localize the flare. We observed the nebula twice, 49 and 55 days after the onset of the gamma-ray flare. We found that there were no rapid changes in the radio emission (at 5.5 GHz) from the nebula which could be attributed to the gamma-ray flare over our 6-day interval. The extrema in the difference image were ∼1.5 mJy beam^{-1}, or 5% of the peak brightness in the nebula, however, we show that most of these difference are in fact artifacts due to deconvolution errors.

The rapid variability of the gamma-ray flares show that the acceleration region must be quite small, with a radius of $\lesssim 10^{-4}$ pc (Bühler & Blandford 2014), corresponding to $\lesssim 0.''01$ at 2 kpc. The acceleration zone for the flares would therefore be completely unresolved in our image. Any features associated with the flare can expand at most at c, it must therefore have a radius of $\lesssim$55 light-days, corresponding to $4.8''$ by the time of our August 26 observations. No such feature is seen in our difference images. Therefore we can say that any radio feature associated with the gamma-ray flare is changing in surface brightness by $\lesssim$0.2 mJy arcsec^{-2} per day at 5 GHz.

We compare our radio images also to an older one from 2001 April, when there was likely no gamma-ray flare. Since the Fermi satellite was launched only on 2008, no gamma-ray observations were available, but since flares seem to have a lifetime on the order of a week and occur on the order of twice per year, the chances that the Crab was flaring during the 2001 April observations are a few percent. Again, no particular compact feature, where the radio emission in a region $< 9.''5$ in diameter has brightened, stands out, although we note that there are changes throughout the nebula. The extremum in the 2012 − 2001 difference image is at 2 mJy arcsec^{-2}, so we can say that any feature associated with the flare had a 5-GHz brightness of less than 2 mJy arcsec^{-2}, or a total of < 145 mJy (assuming a maximal radius of 55 light-days), corresponding to a spectral luminosity of 7×10^{20} erg s^{-1} Hz^{-1}.

We now discuss the radio emission from specific regions

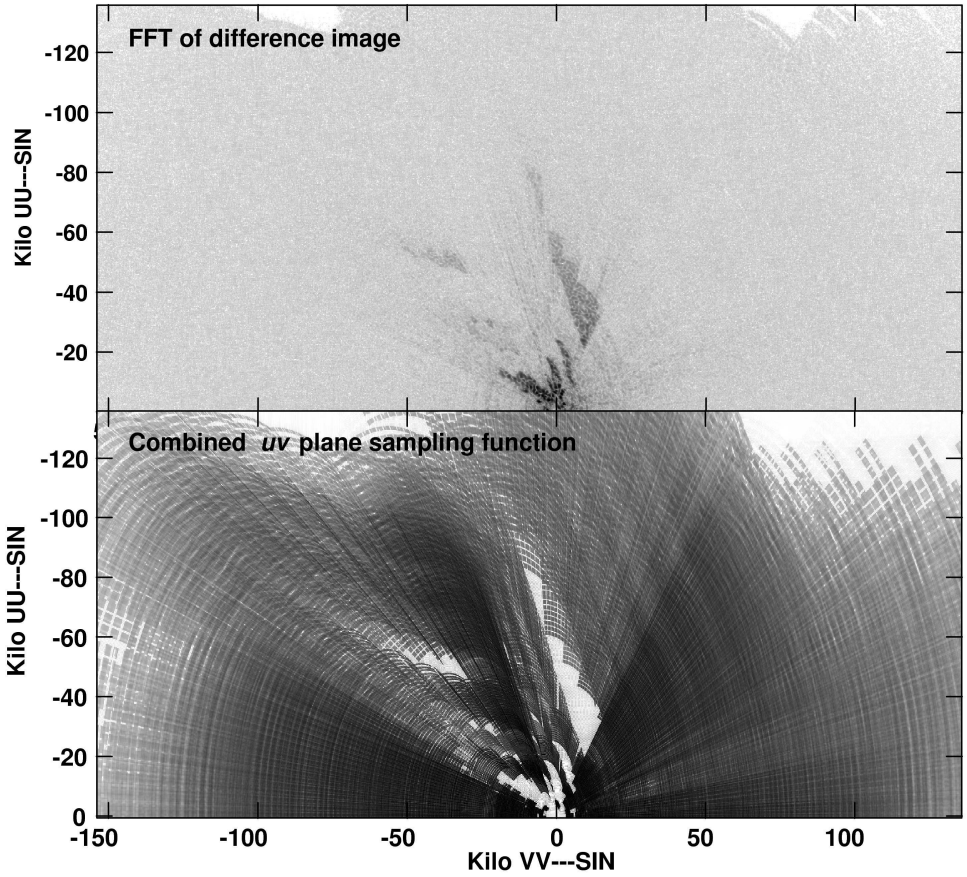

Figure 5. Top: The central region of the Fourier transform (FFT) of the difference image in Figure 3, which was made by subtracting the 2012 August 20 image from the 2012 August 26 one. We show the magnitude of the complex-valued Fourier transform of the image, with 0 being white and positive values being dark. This shows the regions of the u-v plane contributing the most power to the difference image. As the Fourier transform of the real-valued image is symmetrical, we show only half the u-v plane. **B**ottom: The same central region of the FFT of the *product* of the dirty beams corresponding to the two images above. The Fourier transform of the dirty beam is essentially the sampling function in the u-v plane, and the product therefore represents the combined u-v plane sampling function. The blank areas in the product of the two sampling functions, then, show the regions of the u-v plane where we did *not* have measurements at both epochs, and therefore the spatial frequencies where our difference image is poorly constrained by the observations. Note that the *un-sampled* regions correspond very well with those contributing most of the power to the difference image in the top panel.

in the nebula, and whether or not they might represent a radio counterpart of the gamma-ray flares.

6.1 Radio Emission from the Pulsar

We found an unresolved, likely variable, source with flux density $\sim$1 mJy in our images which was coincident with the pulsar. Is this just the pulsed radio emission from the pulsar, averaged over the pulse (period = 33 ms)?

Unfortunately the pulse-averaged flux density of the pulsar is not well known. Moffett & Hankins (1996) found that the pulsar's flux density at 5 GHz is quite variable, and give a mean value of 0.57 mJy for 4.9 GHz. Cordes et al. (2004) also find that above 3 GHz it is quite variable on timescales as short as five minutes. We examined the four 4.8-GHz VLA images of Bietenholz et al. (2004), taken over a period of two months in 2001, and found that the pulsar cannot be reliably identified in them, suggesting that its flux density is < 2 mJy. However, the brightness at the location of the pulsar varies with 0.4 mJy rms over the four images, suggesting that the pulsar's flux density is variable by that amount over times scales of a few weeks. The pulsar was also detected in two 8.4 GHz VLA images with sub-arcsecond resolution taken on 1991 July 14 and Oct. 30 (D. A. Frail and M. F. Bietenholz, unpublished), and had flux densities of $\sim$0.7 and $\sim$0.1 mJy, respectively. Lobanov, Horns & Muxlow (2011) give an upper limit of 0.4 mJy to the pulsar's emission at 4.9 GHz in 2010 Nov. 23. Although it is unknown whether or not a gamma-ray flare had occurred prior to these earlier observations, it is unlikely since flares seem to occur once or twice per year. Taken together these results suggest that the pulse-averaged emission from the pulsar is quite variable, with a mean value at 5 GHz in the range of 0.5 to 1 mJy, and an rms variability of at least 0.4 mJy, with timescales possibly as short as a few minutes.

The flux densities for the central point source we observed in 2012 are therefore within the range of the normal radio emission from the pulsar. Since the point source is coincident with the pulsar to better than $0\farcs14$ (3σ), corre-

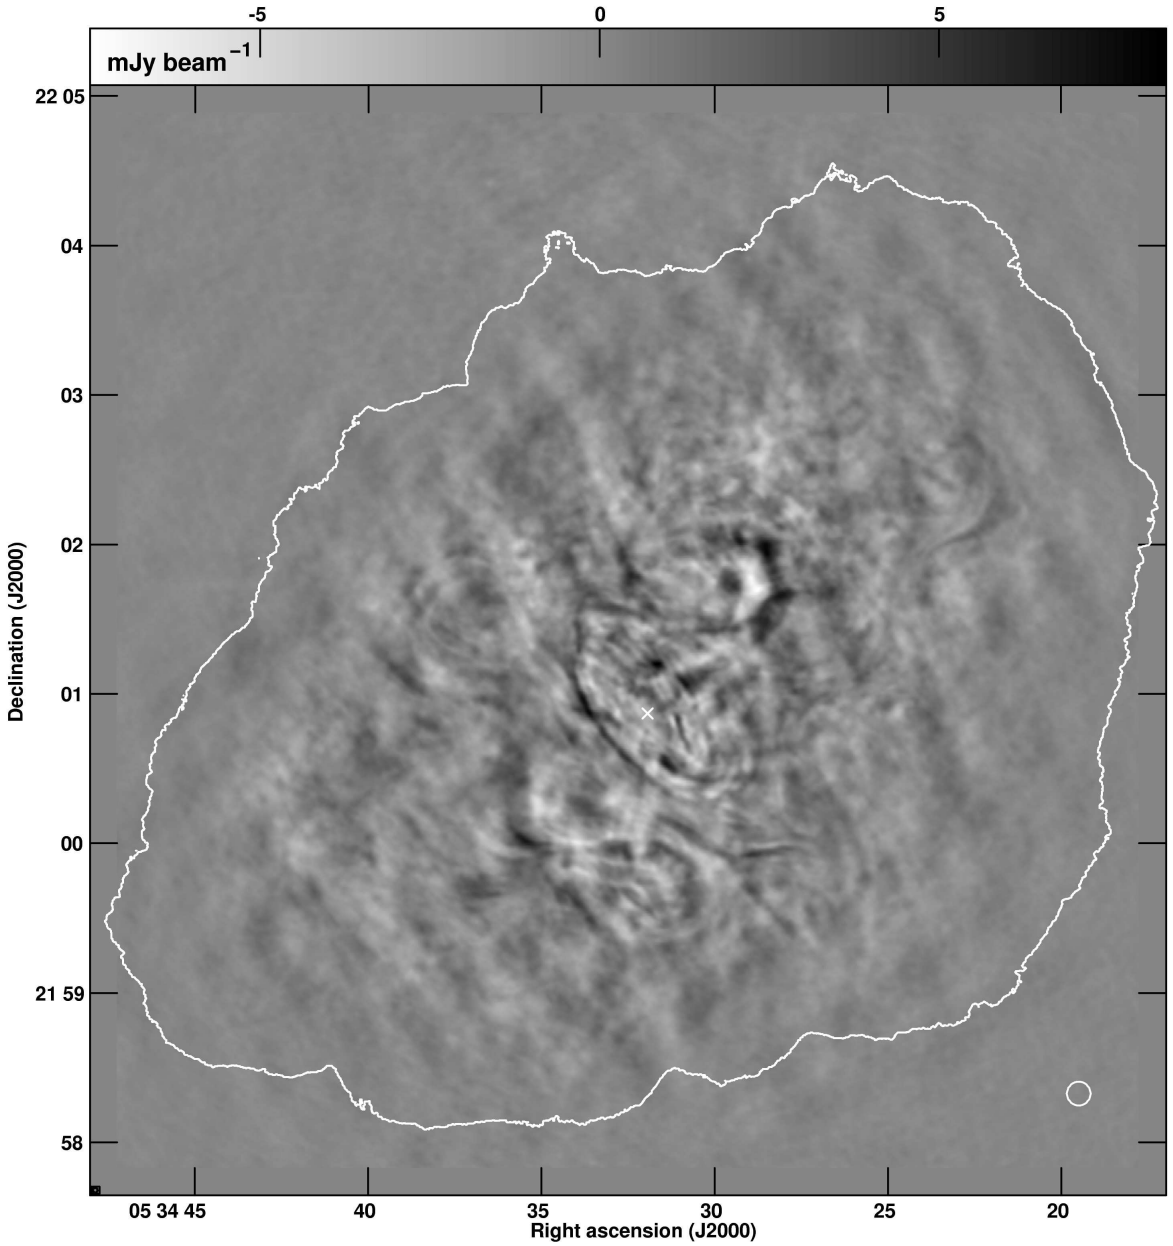

Figure 6. A difference image of the Crab, showing the changes between 2001 April 16 and 2012 August 26 (with the former subtracted from the latter). The difference image has been high-pass filtered at 20″ to isolate structures well sampled by the B-array which was used for both sets of observations. Both were convolved to the same FWHM resolution of 2.00″ × 1.80″ at p.a. 80°. (Note that the greyscale range is much larger than in Fig. 3). As in Fig. 3, the white cross shows the position of the pulsar, the 1% contour of the 2012 image is indicated, and the circle at lower right shows a diameter of 9.″5, which is the diameter that a feature expanding isotropically with speed c since the gamma-ray flare onset on 2012 July 3 would have in 2012 August 26. Before subtraction, the 2001 image was expanded by a factor of 1.016 to account for the expansion of the nebula (see text for details).

sponding to 4.2×10^{15} cm or $\sim 10^{7.4}$ light-cylinder radii, the most parsimonious explanation of the unresolved point source in our images is that it is the pulse-averaged emission from the pulsar itself, and is not directly connected to the gamma-ray flare.

We note that Weisskopf et al. (2013) report a detection of emission coincident with the pulsar at up to 5.0 mJy at 7.8 GHz[4] and 1.78 mJy at 4.9 GHz. The unusually high value of 5.0 mJy was observed only ∼17 days after the onset of the gamma-ray flare. It is large, but not impossibly so for being the normal emission from the pulsar.

[4] Averaged between the values at 7.762 and 7.872 GHz reported in Weisskopf et al. (2013).

6.2 The Inner Knot

There is a compact, variable optical and infra-red feature close to the pulsar called the inner knot, located $\sim$ 0.″65 from the pulsar at p.a. 120° (Hester et al. 1995; Sandberg & Sollerman 2009). It is variable in brightness(Sandberg & Sollerman 2009), and using Hubble Space Telescope images, Sollerman (2003) found the position of this feature to also be variable at the 0.″1 level. None of the reported optical or infra-red positions of the inner knot, however, are closer than 0.″6 to the pulsar. The inner knot is somewhat extended, with Hester et al. (1995) reporting an angular size of 0.″16 in the direction to the pulsar and about three times larger in the perpendicular direction.

Komissarov & Lyutikov (2011) suggested that the inner knot might be the site of the gamma-ray flaring activity on the basis of its proximity to the pulsar and the short synchrotron lifetimes of the gamma-ray emitting electrons.

The central radio point source discussed in § 3.2 above is therefore not associated with the inner knot since the 3σ upper limit on the displacement of the central radio point source from the pulsar is only 0.″14.

We estimated an upper limit on flux density of the inner knot by fitting and subtracting the pulsar from the image shown in Figure 2. From this residual image we estimate that the largest feature which might escape detection at the location of the inner knot is 0.3 mJy beam^{-1}. We therefore find an upper limit to the 5.5-GHz radio emission from the inner knot of 0.3 mJy beam^{-1}. If the inner knot is, as suggested by (Komissarov & Lyutikov 2011), the site of the gamma-ray flaring activity, there is very little corresponding radio activity.

In the infrared, Weisskopf et al. (2013); Sandberg & Sollerman (2009); Melatos et al. (2005) and Sollerman (2003) all report flux densities for the inner knot. The average of the inner knot flux densities reported in the infrared K band (1.4×10^{14} Hz) from those authors is 0.36 mJy. Although Weisskopf et al. (2013) report that the K′ infrared flux rose by 35% following a strong flare, this rise is compatible with the normal variation of the inner knot, and therefore likely not associated with the flare. In the radio we did not detect the inner knot, and therefore also find no evidence that the inner knot shows any response to gamma-ray flares.

Our upper limit on the 5.5-GHz flux density of the inner knot along with the infrared values implies a radio-infrared spectral index, $\alpha^{\rm IR}_{\rm Rad} > +0.02$. By contrast, in the infrared, Melatos et al. (2005) and Sollerman (2003) report that the inner knot's spectral index is ~ -0.8 while Sandberg & Sollerman (2009) suggest a value of -1. The spectral energy distribution of the inner knot therefore must have a turnover somewhere below 10^{14} Hz. The knot's $\alpha^{\rm IR}_{\rm Rad}$ must be rather flatter than the integrated value for nebula, which is ~ -0.4 (e.g. Aharonian & Atoyan 1998; Arendt et al. 2011).

There is another constraint that can be drawn from our present observations and those of Lobanov, Horns & Muxlow (2011) and this is from the limit that can be placed on the flux from a point source at the position of the inner knot. The inner knot has been associated with the region where the pulsar outflow is deflected towards us by a standing termination shock (e.g. Komissarov & Lyutikov 2011). The shock should be located where the momentum flux in the wind, thought to scale as the cosine of the latitude to the fourth power (Tchekhovskoy, Spitkovsky & Li 2013), matches the ambient nebula pressure. If the momentum flux is scaled to the inner edge of the equatorial torus seen in X-rays then the shock radius along the line of sight can be estimated as $\sim 9''$ in angular measure. In this case the deflection angle is $\sim$0.07 rad and the dissipation would have to be substantial. The 5.5 GHz flux density limit on the inner knot is $\sim$ 300 μJy, which implies νL_ν per steradian at 5.5 GHz of 6×10^{26} erg s^{-1} sr^{-1} or 1.6×10^{-11} of the spindown luminosity of the pulsar per steradian.

This surprisingly low radio radiative efficiency of $\lesssim 1.6 \times 10^{-11}$ is a strong constraint on models of the termination shock which should produce some radio synchrotron emission from $\sim$ GeV electrons as well as the low harmonics radiated by more energetic particles. In addition shocks can create coherent radio emission associated with the surface current at the shock front and this, too, will be limited.

6.3 Other Radio Counterparts to Gamma-Ray Flares

Our observations did not detect any radio emission which could be identified as a radio counterpart to the gamma-ray flare. Earlier attempts to find such radio emission by Lobanov, Horns & Muxlow (2011) and Weisskopf et al. (2013) had also proved inconclusive. Unlike some of the earlier observations, we are sensitive to emission anywhere in the body of the nebula, not just near the pulsar. Our observations are also more sensitive to somewhat more extended emission at late times, as might occur if the source were expanding relativistically.

The 1.6-GHz VLBI observations carried out 44 d after a gamma-ray flare in 2011 September did result in a possible detection of radio features associated with that gamma-ray flare: Lobanov, Horns & Muxlow (2011) tentatively identified two compact knots of radio emission, which they termed C1 and C2, as possible sites of the flaring activity. We indicate the two positions of C1 and C2 on our Figs. 2 and 4. No particular feature or change is visible at the position of either of these knots. Lobanov, Horns & Muxlow (2011) report 1.6-GHz flux densities of 0.5 ± 0.3 and 0.4 ± 0.2 mJy for the two knots respectively. On our 5.5-GHz image, no corresponding features can be seen down to a level of $\sim$0.4 mJy. We therefore cannot confirm the radio emission reported by Lobanov, Horns & Muxlow (2011). Weisskopf et al. (2013) also report on radio observations following the strong flare of 2011 April, and report no radio emission features at the locations of C1 and C2 with the best 3σ upper limits on the order of 0.5 mJy. They did not detect any significant variability in the X-ray emission at these locations either.

More generally, we can say that the 5.5-GHz spectral luminosity of any feature associated with the flare in our observations was $\lesssim 2 \times 10^{-4}$ that of the nebula. By contrast, even a relatively weak gamma-ray flare like the one of 2012 July involves an increase of factors of several of the nebular flux above 100 MeV. In other words whatever the process involved in the flares is, they are faint in the radio (this conclusion was already reached in e.g. Weisskopf et al. 2013).

The total energy of the 2012 July flare was of order

10^{41} erg. Given an average magnetic field in the nebula of 300 μG, this energy is the equivalent of the magnetic energy contained in a volume of radius $\sim 2 \times 10^{16}$ cm (0.006 pc).

The energy released in the gamma-ray flares is substantial, with the flare fluxes being up to 1% of the pulsar's spin-down energy loss rate, and the flare timescales implying that this energy release happens in very small regions. It seems likely that the dynamical aftermath of such an energy release would affect the ambient emission, for example through compression or rarefaction of the magnetic field and/or the particle density (Weisskopf et al. 2013). Such disturbances would propagate through the synchrotron plasma filling the bulk of the nebula, likely at the sound speed of $\sim c/\sqrt{3}$. The temporal changes that we observed may therefore be the ripple-like disturbances created in the synchrotron-emitting fluid by the sudden energy releases which produce the gamma-ray flares, propagating through the nebula and interacting with the complex structure of the dense thermal filaments and/or the outer boundary of the nebula. Alternatively, the temporal changes may similarly be the propagating, wave-like effects of instabilities at the termination shock, which instabilities manifest themselves in the rapid variability of the wisps. The rate of rotational energy loss by the pulsar is prodigious at $\sim 5 \times 10^{38}$ erg s^{-1} (e.g. Hester 2008), and most of this energy is thought to be injected into the nebula in the form of magnetic energy (e.g. Porth, Komissarov & Keppens 2014; Bühler & Blandford 2014; Rees & Gunn 1974). Since the majority of this magnetic energy must be dissipated on timescales less than the ∼1000-yr age of the nebula, variation in emissivity on scales from hours to decades is probably not surprising.

The radio-emitting electrons should evolve adiabatically as their synchrotron cooling times are very long compared with the timescales for variation. In this case the synchrotron emissivity at frequency ν varies $\propto B^{5/2-2\alpha}\nu^{\alpha} \propto B^{3.04}\nu^{-0.27}$, where B is the magnetic flux density and we set the radio spectral index $\alpha = -0.27$ (e.g. Bietenholz et al. 1997, and references therein). Suppose first that the γ-ray emission from a large flare is due to an electromagnetic interchange that releases energy $E \sim 10^{42} E_{42}$ erg. The ambient energy density in the central regions of the nebula is $\sim 10^{-7}$ erg cm^{-3} and so we might expect to obtain $O(1)$ changes in the magnetic field and the synchrotron emissivity over a region of size $\ell \sim 10^{16} E_{42}^{1/3}$ cm $\sim 0.5\, E_{42}^{1/3}$ arcsec $\sim 5\, E_{42}^{1/3}$ light days. If we take the effective line-of-sight depth of the nebula to be equivalent[5] to 200″ and the central diffuse surface brightness to be ∼15 mJy arcsec^{-2} (see Fig. 1, 2), then the expected surface brightness change on scale ℓ is $\sim 150\,(\ell/1\mathrm{arcsec})\,\mu$Jy arcsec^{-2}, or $\sim 75\, E_{42}\,\mu$Jy arcsec^{-2}. Since our beam area is ∼1 arcsec^{-2} and is comparable to or larger than the expected values of $\pi\ell^2$, we would in fact only expect changes smaller than $\sim 75\, E_{42}\,\mu$Jy beam^{-1} to appear in our images. Furthermore, for the particular flare of 2012 July, E_{42} was only of order 0.1, so the expected changes in the radio brightness due to the flare would be too small to be detected, especially against the background of radio variability in the nebula that is not directly connected with the flare.

[5] The FWHM of the nebula in the sky plane is ∼200″ at 5.5 GHz, and we take this value as an estimate of its line-of-sight depth.

These estimates of the expected change in the radio emission due to a flare are quite model-dependent and do not take account of beaming. Although the limits on both the surface brightness and the flux discussed above are therefore not, in practice, very constraining, they do support the view that the gamma ray flares have a fairly high radiative efficiency.

6.4 Filamentary structure and temporal variations in the Nebula

Our radio image in Figure 1 shows prominent filamentary structure visible throughout the nebula. Such structure has been seen in earlier radio images (e.g. Bietenholz et al. 2004; Bietenholz, Frail & Hester 2001). At our resolution of ∼ 1.″0, most of the filamentary structure seems resolved.

The filamentary structure visible in our 5.5 GHz is generally well correlated with that at other radio wavelengths from 74 MHz (Bietenholz et al. 1997) to 350 GHz (Arendt et al. 2011; Green, Tuffs & Popescu 2004), although the images available at these other frequencies are of lower resolution. At long infrared wavelengths the filamentary structure still largely corresponds to that visible in radio, but at shorter infrared wavelengths (Temim et al. 2006) and in optical (Loll et al. 2013) much of the filamentary structure has disappeared. Since the optical continuum emission is synchrotron emission like the radio, it might be expected to have a similar morphology, but instead shows much less filamentary structure. A possible reason for this difference may just be synchrotron burnoff. Much of the filamentary structure resides in the outer nebula, and if all the relativistic particles originate in the termination shock near the pulsar, then the synchrotron lifetimes of the higher-energy electrons are too short to reach, and therefore illuminate in the optical, the outer filaments (Tang & Chevalier 2012). It has also been proposed, however, that the radio-emitting particles are accelerated in the vicinity of the filaments, which would also be consistent with the observed decrease of the filamentary structure towards the optical.

Furthermore, although we did not find any particular radio variability which might be associated with the gamma-ray flare, we did find that over a period of about a decade there were striking changes in the radio emission occurring throughout the nebula, which were most pronounced within ∼1 pc of the pulsar. Such changes had previously been reported (Bietenholz et al. 2004; Bietenholz, Frail & Hester 2001; Bietenholz & Kronberg 1992), but are now more clearly visible. The temporal variations were up to 10% of the peak brightness of the nebula (at ∼1.″9 resolution). The changes have a complex morphology of arcs and knots: there is little structure at the ∼2″ level, but considerable structure on spatial scales of a few arcsec (∼0.04 pc) up to perhaps 1′ (∼0.6 pc). Earlier, Bietenholz et al. (2004) had examined a differences in radio image made ∼3 yr apart, and found similar changes, but which were mostly confined to a region within ∼1′ from the pulsar. Over our longer timescale of about a decade there are significant changes at least 2′ from the pulsar. This suggests that the variability timescale becomes longer at larger distances from the pulsar.

We note that exactly such a pattern of radio brightness changes is seen in the dynamical relativistic magneto-hydrodynamic (MHD) model of Olmi et al. (2014). Their

model was axisymmetric and had parameters chosen to best reproduce the Crab's X-ray morphology, but unlike most previous MHD modelling, Olmi et al. examined the morphology of the resulting radio rather than the high-energy emission. They find that the radio morphology is essentially the same regardless of whether the radio-emitting electrons are continuously accelerated at the termination shock or if they are uniformly distributed in the nebula. The radio morphology (including the radio wisps) therefore reflects only the underlying flow structure, but not the site where the radio-emitting electrons are accelerated. In their model the injected magnetic field from the pulsar changes polarity around the rotational equator, and eddies then cause the current sheet to twist and tangle downstream of the termination shock, and it is these instabilities which give rise to the observed variability in the radio brightness, and produces synthesized radio difference images very much like the observed one in Figure 6. In agreement with our observations the modelled radio brightness changes in the inner nebula are more rapid and of larger amplitude than those in the outer nebula. Our observations therefore lend considerable credence to the MHD modelling of Olmi et al. (2014).

7 CONCLUSIONS AND SUMMARY

We briefly summarize our results and conclusions:

(1) We present new deep 5.5-GHz radio images of the Crab Nebula, about two months after a gamma-ray flare event in 2012 July.
(2) We find no significant change in the Crab's radio emission localized to a region < 2 light-months in radius. Any radio counterpart to the gamma-ray flare has a total flux density of < 145 mJy, corresponding to a spectral luminosity of $< 7 \times 10^{20}$ erg s^{-1} Hz^{-1}, or $< 2 \times 10^{-4}$ that of the nebula. The surface brightness of any flare counterpart is < 2 mJy arcsec^{-2}.
(3) The low limits for a radio counterpart to the flare imply that few low-energy electrons ($\lesssim$10 GeV) are accelerated in the flaring event. Nonetheless, the energy release of the flare is expected to produce changes in the magnetic field, which should produce corresponding changes in the radio brightness. We show however, that such changes are likely to be smaller than our observed radio brightness limits on energetic grounds.
(4) We detect radio emission at the mJy level from within $0\farcs14$ (4.2×10^{15} cm) of the pulsar. This emission is very likely just the normal pulse-averaged emission from the pulsar, but we cannot rule out mJy-level radio emission associated with the gamma-ray flare from very near the pulsar at this level.
(5) We find no discernible radio emission from the "inner knot", seen at $\sim 0\farcs65$ from the pulsar in the optical and infra-red. We set an upper limit of 0.3 mJy beam^{-1} on the 5.5-GHz radio brightness of the inner knot. This limit represents a very low radiative efficiency of 1.6×10^{-11} of the spindown luminosity of the pulsar per steradian.
(6) We find that deconvolution errors are several times larger than the thermal noise, even in these images made using the wide bandwidth and consequently excellent u-v coverage of the VLA. The deconvolution errors represent a fraction of 1.2% of the peak brightness of the nebula, although this fraction is likely strongly dependent on the u-v coverage and source geometry.
(7) By comparing our images to ones from 2001, we find widespread changes in the brightness over decade timescales, as large as 2 mJy arcsec^{-2}, or up to $\sim$10% of the peak brightness of the nebula. These changes are both larger in amplitude and morphologically distinct from the deconvolution errors. Averaged over the nebula the changes in surface brightness over decade timescales have an rms of 1.4% of the peak brightness. (These changes are in addition to the secular decay in brightness of $\sim$1.3% per decade). The morphology of the changes is complex, suggesting both filamentary and knotty structures. The variability is stronger in the centre of the nebula, and the timescales are likely shorter near the centre than at the periphery. These variations correspond well to those seen in MHD simulations of the Crab by (Olmi et al. 2014).

ACKNOWLEDGEMENTS

We are indebted to the late Dr. Michael Gaylard for his encouragement and support of astronomical research at Hartebeesthoek Radio Astronomy Observatory. We also thank David T. Thompson for comments on the manuscript, and Jon Arons, Martin Weisskopf and Claire Max for support and encouragement. We thank the anonymous referee for useful comments on the manuscript. Research at Hartebeesthoek was partly supported by National Research Foundation (NRF) of South Africa. Research at York University was partly supported by NSERC. R. Blandford acknowledges support from NSF grant AST 1212195. The National Radio Astronomy Observatory is a facility of the National Science Foundation operated under cooperative agreement by Associated Universities, Inc. We have made use of NASA's Astrophysics Data System Bibliographic Services.

REFERENCES

Abdo A. A. et al., 2011, Science, 331, 739
Aharonian F. A., Atoyan A. M., 1998, in Neutron Stars and Pulsars: Thirty Years after the Discovery, p. 439
Aller H. D., Reynolds S. P., 1985, ApJL, 293, L73
Arendt R. G. et al., 2011, ApJ, 734, 54
Arons J., 2012, SSRv, 173, 341
Bietenholz M. F., Frail D. A., Hester J. J., 2001, ApJ, 560, 254
Bietenholz M. F., Hester J. J., Frail D. A., Bartel N., 2004, ApJ, 615, 794
Bietenholz M. F., Kassim N., Frail D. A., Perley R. A., Erickson W. C., Hajian A. R., 1997, ApJ, 490, 291
Bietenholz M. F., Kronberg P. P., 1990, ApJL, 357, L13
Bietenholz M. F., Kronberg P. P., 1991, ApJ, 368, 231
Bietenholz M. F., Kronberg P. P., 1992, ApJ, 393, 206
Bietenholz M. F., Kronberg P. P., Hogg D. E., Wilson A. S., 1991, ApJL, 373, L59
Buehler R. et al., 2012, ApJ, 749, 26
Bühler R., Blandford R., 2014, Reports on Progress in Physics, 77, 066901
Bykov A. M., Pavlov G. G., Artemyev A. V., Uvarov Y. A., 2012, MNRAS, 421, L67

Cerutti B., Werner G. R., Uzdensky D. A., Begelman M. C., 2013, ApJ, 770, 147
Clausen-Brown E., Lyutikov M., 2012, MNRAS, 426, 1374
Cordes J. M., Bhat N. D. R., Hankins T. H., McLaughlin M. A., Kern J., 2004, ApJ, 612, 375
Espinoza C. M., Jordan C., Stappers B. W., Lyne A. G., Weltevrede P., Cognard I., Theureau G., 2010, The Astronomer's Telegram, 2889, 1
Green D. A., Tuffs R. J., Popescu C. C., 2004, MNRAS, 355, 1315
Hester J. J., 2008, ARA&A., 46, 127
Hester J. J. et al., 1995, ApJ, 448, 240
Hibschman J. A., Arons J., 2001, ApJ, 560, 871
Hogg D. E., MacDonald G. H., Conway R. G., Wade C. M., 1969, AJ, 74, 1206
Komissarov S. S., 2013, MNRAS, 428, 2459
Komissarov S. S., Lyutikov M., 2011, MNRAS, 414, 2017
Lobanov A. P., Horns D., Muxlow T. W. B., 2011, A&A, 533, A10
Loll A. M., Desch S. J., Scowen P. A., Foy J. P., 2013, ApJ, 765, 152
Lyutikov M., 2014, Astronomische Nachrichten, 335, 227
Lyutikov M., Balsara D., Matthews C., 2012, MNRAS, 422, 3118
Melatos A. et al., 2005, ApJ, 633, 931
Meyer M., Horns D., Zechlin H.-S., 2010, A&A, 523, A2
Mickaliger M. B. et al., 2012, ApJ, 760, 64
Moffett D. A., Hankins T. H., 1996, ApJ, 468, 779
Narayan R., Nityananda R., 1986, ARA&A., 24, 127
Nolan P. L. et al., 2012, ApJS, 199, 31
Ojha R., Buehler R., Hays E., Dutka M., 2012, The Astronomer's Telegram, 4239, 1
Olmi B., Del Zanna L., Amato E., Bandiera R., Bucciantini N., 2014, MNRAS, 438, 1518
Porth O., Komissarov S. S., Keppens R., 2014, MNRAS, 438, 278
Rees M. J., Gunn J. E., 1974, MNRAS, 167, 1
Sandberg A., Sollerman J., 2009, A&A, 504, 525
Sault R. J., Oosterloo T. A., 1996, Imaging algorithms in radio interferometry, Stone W. R., ed., Oxford University Press (Oxford), pp. 30; arXiv:astro–ph/0701171
Sollerman J., 2003, A&A, 406, 639
Striani E. et al., 2011, ApJL, 741, L5
Sturrock P., Aschwanden M. J., 2012, ApJL, 751, L32
Tang X., Chevalier R. A., 2012, ApJ, 752, 83
Tavani M. et al., 2011, Science, 331, 736
Tchekhovskoy A., Spitkovsky A., Li J. G., 2013, MNRAS, 435, L1
Temim T. et al., 2006, AJ, 132, 1610
Vinyaikin E. N., 2007, Astronomy Reports, 51, 570
Weiland J. L. et al., 2011, ApJS, 192, 19
Weisskopf M. C. et al., 2013, ApJ, 765, 56
Wilson A. S., 1972, MNRAS, 160, 373